

26
CLAIMS

1. An InfiniBand architecture subnet, comprising:
a master subnet manager function;
5 database elements of the InfiniBand architecture subnet, wherein the master subnet manager function updates the database elements;
a replicated set of the database elements; and
a set of standby subnet managers, wherein the replicated set of the database elements is created at each of the set of standby subnet managers, wherein a standby
10 subnet manager included in the set of standby subnet managers assumes the master subnet manager function, and wherein the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function initializes the InfiniBand architecture subnet using the replicated set of the database elements.
- 15 2. The InfiniBand architecture subnet of claim 1, wherein the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function reinitializes the InfiniBand architecture subnet.
- 20 3. The InfiniBand architecture subnet of claim 1, wherein the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function manages the InfiniBand architecture subnet using the replicated set of the database elements.
- 25 4. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements is created at each of the set of standby subnet managers out of band of the InfiniBand architecture subnet.
- 30 5. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements is created at each of the set of standby subnet managers inband on the InfiniBand architecture subnet.

6. The InfiniBand architecture subnet of claim 5, wherein the replicated set of the database elements is created at each of the set of standby subnet managers using reliable multi-packet transaction protocol.

5 7. The InfiniBand architecture subnet of claim 5, wherein the replicated set of the database elements is created at each of the set of standby subnet managers using reliable connection transport service.

10 8. The InfiniBand architecture subnet of claim 5, wherein the replicated set of the database elements is created at each of the set of standby subnet managers using reliable datagram transport service.

15 9. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements comprises an event subscription.

10 10. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements comprises a multicast record.

20 11. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements comprises a service record.

25 12. The InfiniBand architecture subnet of claim 11, wherein the service record comprises a lease time, wherein the master subnet manager function converts the lease time to a first end time, wherein the master subnet manager function converts the first end time to a remaining time, wherein the standby subnet manager included in the set of standby subnet managers converts the remaining time to a second end time, and wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

30 13. The InfiniBand architecture subnet of claim 11, wherein the master subnet manager function periodically decrements a lease time, wherein the lease time becomes a remaining time, wherein the standby subnet manager included in the set of standby subnet

managers converts the remaining time to a second end time, and wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

5 14. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements comprises an extended node record.

10 15. The InfiniBand architecture subnet of claim 1, wherein the set of standby subnet managers is selected based on a priority value and a globally unique identifier.

16. The InfiniBand architecture subnet of claim 1, wherein the replicated set of the database elements are periodically updated.

15 17. An InfiniBand architecture node, comprising:
a master subnet manager function; and
database elements of an InfiniBand architecture subnet, wherein the master subnet manager function updates the database elements, wherein the master subnet manager function initiates creation of a replicated set of database elements at a set of standby subnet managers, wherein a standby subnet manager included in the set of standby subnet managers assumes the master subnet manager function, and wherein the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function initializes the InfiniBand architecture subnet using the replicated set of the database elements.

20 18. The InfiniBand architecture node of claim 17, wherein the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function reinitializes the InfiniBand architecture subnet.

25 19. The InfiniBand architecture node of claim 17, wherein the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function manages the InfiniBand architecture subnet using the replicated set of the database elements.

20. The InfiniBand architecture node of claim 17, wherein the replicated set of the database elements is created at each of the set of standby subnet managers out of band of the InfiniBand architecture subnet.

5

21. The InfiniBand architecture subnet of claim 17, wherein the replicated set of the database elements is created at each of the set of standby subnet managers inband on the InfiniBand architecture subnet.

10

22. The InfiniBand architecture node of claim 21, wherein the replicated set of the database elements is created at each of the set of standby subnet managers using reliable multi-packet transaction protocol.

15

23. The InfiniBand architecture node of claim 21, wherein the replicated set of the database elements is created at each of the set of standby subnet managers using reliable connection transport service.

20

24. The InfiniBand architecture node of claim 21, wherein the replicated set of the database elements is created at each of the set of standby subnet managers using reliable datagram transport service.

25. The InfiniBand architecture node of claim 17, wherein the replicated set of the database elements comprises an event subscription.

25

26. The InfiniBand architecture node of claim 17, wherein the replicated set of the database elements comprises a multicast record.

30

27. The InfiniBand architecture node of claim 17, wherein the replicated set of the database elements comprises a service record.

28. The InfiniBand architecture node of claim 27, wherein the service record comprises a lease time, wherein the master subnet manager function converts the lease

time to a first end time, wherein the master subnet manager function converts the first end time to a remaining time, wherein the standby subnet manager included in the set of standby subnet managers converts the remaining time to a second end time, and wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

29. The InfiniBand architecture node of claim 27, wherein the master subnet manager function periodically decrements a lease time, wherein the lease time becomes a remaining time, wherein the standby subnet manager included in the set of standby subnet managers converts the remaining time to a second end time, and wherein the second end time is a function of the remaining time and a local time at the standby subnet manager included in the set of standby subnet managers.

30. The InfiniBand architecture node of claim 17, wherein the replicated set of the database elements comprises an extended node record.

31. An InfiniBand architecture node comprising a computer-readable medium containing computer instructions for instructing a processor to perform a method of replicating database elements in an InfiniBand architecture subnet, the instructions comprising:

- a master subnet manager function updating database elements of the InfiniBand architecture subnet;
- creating a replicated set of the database elements at each of a set of standby subnet managers;
- a standby subnet manager included in the set of standby subnet managers assuming the master subnet manager function; and
- the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function initializing the InfiniBand architecture subnet using the replicated set of the database elements.

30

32. The InfiniBand architecture node of claim 31, wherein initializing comprises reinitializing the InfiniBand architecture subnet.

33. The InfiniBand architecture node of claim 31, further comprising the standby subnet manager included in the set of standby subnet managers that assumes the master subnet manager function managing the InfiniBand architecture subnet using the replicated set of the database elements.

34. The InfiniBand architecture node of claim 31, wherein creating comprises creating a replicated set of the database elements at each of a set of standby subnet managers out of band of the InfiniBand architecture subnet.

35. The InfiniBand architecture node of claim 31, wherein creating comprises creating a replicated set of the database elements at each of a set of standby subnet managers inband on the InfiniBand architecture subnet.

36. The InfiniBand architecture node of claim 35, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable multi-packet transaction protocol.

37. The InfiniBand architecture node of claim 35, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable connection transport service.

38. The InfiniBand architecture node of claim 35, wherein creating inband on the InfiniBand architecture subnet comprises creating using reliable datagram transport service.

39. The InfiniBand architecture node of claim 31, wherein the replicated set of the database elements comprises an event subscription.

40. The InfiniBand architecture node of claim 31, wherein the replicated set of the database elements comprises a multicast record.

41. The InfiniBand architecture node of claim 31, wherein the replicated set of the database elements comprises a service record.

5 42. The InfiniBand architecture node of claim 41, further comprising:
the service record comprising a lease time;
the master subnet manager function converting the lease time to a first end time;
the master subnet manager function converting the first end time to a remaining
time; and

10 the standby subnet manager included in the set of standby subnet managers
converting the remaining time to a second end time, wherein the second end time is a
function of the remaining time and a local time at the standby subnet manager included in
the set of standby subnet managers.

15 43. The InfiniBand architecture node of claim 41 further comprising:
the master subnet manager function periodically decrementing a lease time;
the lease time becoming a remaining time; and
the standby subnet manager included in the set of standby subnet managers
converting the remaining time to a second end time, and wherein the second end time is a
function of the remaining time and a local time at the standby subnet manager included in
20 the set of standby subnet managers.

44. The InfiniBand architecture node of claim 31, wherein the replicated set of the database elements comprises an extended node record.

25 45. The InfiniBand architecture node of claim 31, further comprising selecting the
set of standby subnet managers based on a priority value and a globally unique identifier.

30 46. The InfiniBand architecture node of claim 31, wherein creating the replicated
set of the database elements comprises periodically updating the replicated set of the
database elements.